METHOD PROVIDING DISTRIBUTION MEANS FOR REFERENCE CLOCKS ACROSS PACKETIZED NETWORKS

Cross-Reference to Related Application

[0001] This application claims the benefit under 35 U.S.C. 119(e) of U.S. Provisional Application No. 60/448,738 filed February 20, 2003, the contents of which are incorporated by reference herein.

Field of the Invention

[0002] The present invention relates to the field of digital communications. More specifically, the present invention relates to a method of distributing fine-grained clocks across packet networks.

Background of the Invention

[0003] There is increasing interest in providing telecommunications services over packet networks that were formerly provided over TDM (time division multiplexed) networks. Most of these telecommunications services require a synchronized clock, which is not provided by the physical layer of a packet network in the same way as it was provided by the physical layer of the TDM network. Various approaches to synchronization over packet networks have been proposed and standardized (see NTP and IEEE Std 1588, among others), but these methods usually distribute a single time-of-day reference. This is not a good match for some services which formerly ran over plesiochronous TDM network connections, where each connection was allowed a certain frequency inaccuracy, hence multiple connections entailed multiple reference clocks. For the sake of backwards compatibility, it is desirable for this possibility of multiple reference clocks to be supported over packet networks. A brute-force approach to this issue is to use adaptive clock recovery, which tries to average out packet rates to calculate the bit clock. Adaptive clock recovery tends to have poor wander performance, even when timestamping and packet sequence numbering are used to overcome packet-loss issues.

Summary of the Invention

[0004] Disclosed is a method to distribute a number of reference clocks across a packet

network. The described method uses a single "basis clock" per node; this basis clock is synchronized in each of the slave nodes to a master node. The master node also numerically encodes a number of reference clocks with respect to the master's basis clock. At each slave node, with the slave node's basis clock synchronized to the master basis clock, and with numerical information describing the reference clock(s) with regard to the master's basis clock, the reference clock(s) can be reconstructed with high accuracy.

[0005] Thus, according to one aspect, the invention provides a method of distributing a number of reference clocks across a packet network. The method comprises providing a basis clock in a master node and one or more slave nodes, encoding multiple reference clocks with respect to the basis clock in the master node to generate numerical information describing the reference clock(s) in relation to the basis clock in the master node, synchronizing the basis clock in each of the slave nodes to the basis clock in the master node using time-stamped synchronization packets, and recovering said one or more reference clocks at the slave nodes using said numerical information describing the reference clock(s) in relation to the basis clock in the master node.

[0006] In another aspect, the invention provides a packet network comprising a master node and one or more slave nodes, said master node and said slave nodes having basis clocks, means for numerically encoding a plurality of reference clocks with respect to the basis clock in the master node to generate numerical information describing the reference clock(s) with regard to the master's basis clock, a sender for sending time-stamped synchronization packets to said one or more slave nodes, a receiver at the slave nodes for receiving said time-stamped synchronization packets and synchronizing the basis clocks in the slave nodes with the basis clock in the master node, and means at the slave nodes for recovering said reference clocks using said numerical information describing the reference clock(s) with regard to the master's basis clock.

[0007] There are many advantages in using the disclosed method. The method works on existing networks without proprietary modifications to the hubs, switches and routers. It is possible to efficiently encode multiple reference clocks with the possibility of using minimum bandwidth per reference clock. In addition, it offers better wander performance than Adaptive Clock Recovery methods, when multiple reference clocks

are present. Such a situation may arise when a number of T1 or E1 lines come into the source node, each having its own primary reference clock.

[0008] Other aspects and advantages of embodiments of the invention will be readily apparent to those ordinarily skilled in the art upon a review of the following description.

Brief Description of the Drawings

[0009] Embodiments of the invention will now be described in conjunction with the accompanying drawings, wherein:

Figure 1 illustrates a global picture of clocks distributed in the network; and

Figure 2 illustrates the master DCO versus master local time.

[0010] This invention will now be described in detail with respect to certain specific representative embodiments thereof, the materials, apparatus and process steps being understood as examples that are intended to be illustrative only. In particular, the invention is not intended to be limited to the methods, materials, conditions, process parameters, apparatus and the like specifically recited herein.

Detailed Description of the Preferred Embodiments

[0011] Figure 1 shows a simple configuration with two slave nodes and a single reference clock. The local oscillator which is the basis clock for slave node Q is characterized against the local oscillator which is the basis clock for master node P. These local oscillators offer one place to perform a cost/benefit tradeoff – better, more expensive oscillators provide better performance.

[0012] Figure 1 shows one reference clock "Clock A". A series of alignment calculations is performed.

[0013] Alignment 1: Reference clock versus master DCO

[0014] The DCO in the PPLL block at the master node will be locked to the incoming reference Clock A in a standard manner for digital PLLs: A PI-controller will lock the master DCO (acquisition PLL) to the reference clock.

3

[0015] Alignment 2: Master DCO versus master local time

[0016] The DCO ramp of the master describes a saw tooth as a function of its local time and represents unambiguously both phase and frequency of the reference clock (Figure 2).

[0017] In the P-PLL the DCO phase will be extended with sufficient leading bits to prevent wrapping of the DCO phase too frequently for the software (dotted line). So at the master the local software samples the dotted line. This line is described with the simple equation:

$$\varphi_{DCO_M} = m_0 + m_1 \cdot t_M \qquad (1)$$

[0018] The software at the master will continuously determine the constants $M_{0,1}$ by sufficient sampling and curve fitting. MO indicates the absolute phase of the DCO at local time $t_M = 0$. MI indicates the frequency of the DCO measured relative to the basis clock (local time) of the Master. The two M-constants are unambiguously describing the relationship between the local oscillator basis clock (local time) at the master and the ramp of the master DCO that is locked to the reference clock. These constants are pair-wise transmitted to the slave nodes. The reference clock is thus numerically encoded with respect to the master basis clock.

[0019] Alignment 3: Local time of the master versus local time of the slave

[0020] The local time of the master and the slave are unrelated and will thus have a slight mutual frequency offset. This creates time scales drifting relative to each other. When the oscillators are reasonable steady this drift will not change much. Using broadcasting methods (referenced as broadcast events) or NTP-like protocols, the alignment between the master basis clock and the slave basis clock (the two local time domains) are described with:

$$t_M = c_0 + c_1 \cdot t_S \qquad (2)$$

[0021] These constants are determined by curve fitting at the slave. They

4

unambiguously describe the relationship between master basis clock (local master time) and slave basis clock (local slave time). Any event (e.g. clock edges) at the master can now be translated to the same real time at the slave using above equation.

[0022] Alignment 4: Local time of the slave versus the slave DCO

[0023] By sending the master's M-constants to the slave, the slave can reconstruct master DCO ramp exactly at the slave by substituting the second equation below:

$$\varphi_{DCO_M} = m_0 + m_1 \cdot t_M \qquad (3)$$

$$\Longleftrightarrow$$

$$\varphi_{DCO_M} = \varphi_{DCO_S} = m_0 + m_1 \cdot (c_0 + c_1 \cdot t_S) \qquad (4)$$

[0024] The result is that all individual 4 alignments are unambiguously aligned, which means that the phase of the Master and the Slave clock are aligned (locked) too.

[0025] It can be seen by extension that if multiple acquisition DCOs are provided at the master, then multiple reference clocks can be encoded. The cost of each reference clock is the packet traffic to transport the M-constants pair for that reference clock. The rate at which the M- constants pair should be transmitted is related to the stability of the basis clock, and to the low speed variation of the reference clock with respect to the basis clock.

[0026] There are a number of variations possible within the scope of the disclosed invention. The packet network can be Ethernet (over various physical layer speeds; 10 Mbps, 100 Mbps, 1 Gbps,etc.), a wireless packet network such as 802.1 1, ATM, Frame relay, or any other standard or proprietary packet network. The Packet Network may or may not contain mechanisms for priority and Quality of Service/Class of Service.

[0027] The basis clock can be ovenized crystal oscillator, temperature compensated crystal oscillator, a real-time (time and date) clock, or any other oscillator with accuracy and stability suiting the application.

[0028] The slave basis clock can be synchronized to, or measured against, the master basis clock via NTP (Network Time Protocol), SNTP (Simple Network Time Protocol), PTP (Precise Time Protocol), a timing pulse on the extra Ethernet pair, or any other standard or proprietary methods.

[0029] Transmission of packets containing reference clock information can happen at regular preset intervals, chosen to minimize packet network loading while still providing sufficient bandwidth to track low speed variations in the reference clock, at pseudo-random intervals chosen to avoid any possible repetitious pattern of other packet network trait by "noise-shaping" the transmissions, or at auto-configurable intervals based on information gathered at the acquisition DPLL about the low speed variations in the reference clock.

[0030] Numerous modifications may be made without departing from the spirit and scope of the invention as defined in the appended claims.